SEMI Workshop

Day 2

Possible Working Sessions (1)

1. Trust domain
2. Path Characteristics API
3. Middlebox interaction
4. Middlebox measurements
5. User-space stacks
6. Balance of power (encryption)

Possible Working Sessions (2)

Trust domain

- Incl. incentive structure?

Path Characteristics API

- How to get control of end-2-end in-network functions back to the client?
- Middlebox detection mechanisms

Middlebox interaction

- Which information would be needed?
- Protocol design

Possible Working Sessions (3)

Middlebox measurements

- Failure reports in happy-eyeball protocols
- Define methodology and data model for aggregation?

User-space stacks

- Can this support faster deployment of new protocols/protocol extensions?
- Incl. UDP guidance?

Balance of power

- Using encryption to provide extensibility of existing protocols
- WG charter?

Other work items

- „UDP Encaps for Dummies“ — Non-wg-forming BoF
- Defining „good/bad“ middlebox behavior
- …

Requirements

- Deploy: existing Internet/kernels, user space, not root
- Choices for congestion, retransmit, etc.
- Single firewall-traversal mechanism,
  multiple transport semantics
- Multiple interfaces for each endpoint
- Low overhead
- Determine protocol in use (fast, but not port-specific)
- Associate packets with a flow (fast)
- Policy per-flow
- In-flow path->application (p2a) and application->path (a2p)
  - treat as ignorable hints both ways unless authorized
  - p2a, a2p have separate security context from e2e

Session protocol for UDP Datagrams (SPUD)

- UDP for per-app demux
- Magic number
- Session ID
- Command (open, close, error)
- Transport inside (e.g. TCP, SCTP)
- Path-to-app flag, App-to-path flag
  - ------
- Simple transport for p2a, a2p
  - Frags
  - CBOR (RFC 7049, http://cbor.io/)
  - MUST ignore not-understood

A2P examples

- Request for special handling
  - e.g., I would prefer you drop rather than delay
- Application/device capabilities
- Measurement probes, requests

P2A examples

- MTU
- Bandwidth
- ECN
- Pacing
- Timeout info
- Back-scatter: find non-compliant path elements?
- Errors with URL for reporting, diagnosis
- Non-error status for determining break point (ping++ on path)

API

- Event framework
  - Originally in library
  - Some parts move to kernel if successful
  - Application can send and receive events
  - Good timing info very helpful
- Different layers
  - Debug: bytes sent/received
  - Session: open, close, error
  - Path: CBOR send/receive
  - Transport: specific to semantics (e.g. new SCTP)

Open Source

- Library
- Client, server, path
- Lots of languages
- Implement at least some a2p, p2a
  (or it will never deploy)
- At least one transport

IP-UDP	Header	Encrypted data, Checksum

- Transport innovation requires encryption
- Encryption – and security – is hard
- Let’s reuse TLS!

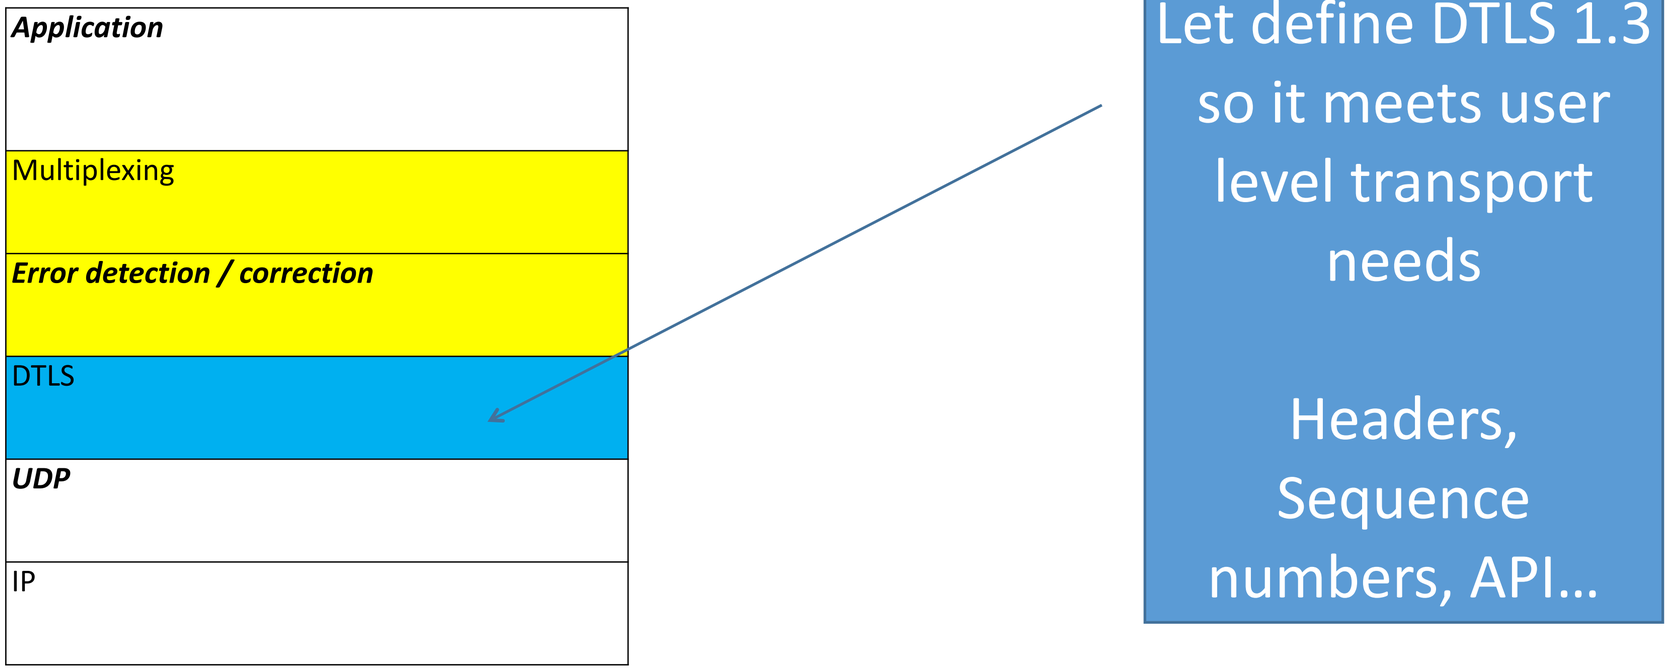

Let define DTLS 1.3 so it meets user level transport needs

Headers, Sequence numbers, API…

SPUD datagram from port A-B, encapsulated in a number of headers

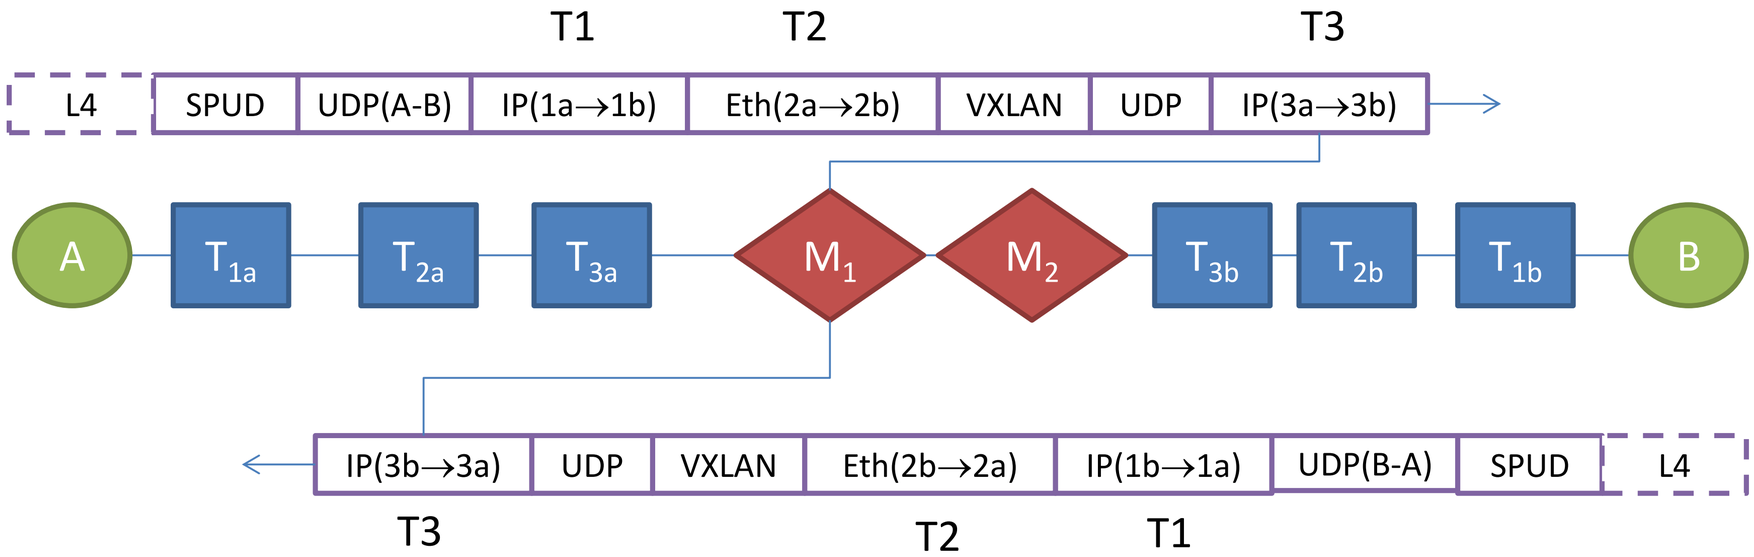

If middlebox 1 wants to create a SPUD datagram from Bto A,
it has to turn round all the src-dest pairs on all the encapsulating headers

Tunnel protocol = any of:

- VXLAN
- GRE
- GTPv1
- GTPv2
- L2TP
- IPv4
- IPv6
- ANother tunnel
- Future tunnel protocols

Action Items (1)

- TPC, all: Minutes for Dallas
- TPC, all: Workshop report
- Brian, Mirja: post slides from workshop on page
- Aaron, Eliot: Organize Dallas Bar BOF on client-side middlebox detection and error information aggregation.
- Ted: WebRTC use case for subtransport session (SPUD) -00 draft
- Joe: Minimum-SPUD -00 draft
- Brian, Eliot, Mirja: Minimum-SPUD Dallas (non-WG?) BoF proposal
- Bob: semi-workshop@iab.org thread to lead to recommendation on IAB statement on basic assumptions about transport evolution.
- Christian, EKR, Jana: DTLS 1.3/1.+ as subtransport session -00:
  - What does DTLS already provide?
  - What could we add to DTLS that would make this easier?

Action Items (2)

- Bob, Mirja, Jana: Organize Dallas Bar BOF + tsvarea preso on transport protocol extensibility
  - define cryptographic protocol based approaches to transport protocol extensibility
- Bob: Review 3234 to see if it should be bis'd
- Gorry, David: figure out how to do UDP Encap for Dummies given 5405bis/ tsvarea +apparea preso
- Marc: publish internetover443
- Joe, Dave: organize apparea preso on the weaponization of HTTPS
- Brian: Plenary presentation summarizing SEMI (Joe provides slides)
- Bob: Presentation summarizing SEMI to ETSI NFV